The Secrecy of Compressed Sensing Measurements

Yaron Rachlin and Dror Baron

Abstract—Results in compressed sensing describe the feasibility of reconstructing sparse signals using a small number of linear measurements. In addition to compressing the signal, do these measurements provide secrecy? This paper considers secrecy in the context of an adversary that does not know the measurement matrix used to encrypt the signal. We demonstrate that compressed sensing-based encryption does not achieve Shannon’s definition of perfect secrecy, but can provide a computational guarantee of secrecy.

I. INTRODUCTION

The theory of compressed sensing, described in [1–3] and surveyed in [4, 5], demonstrates the feasibility of recovering sparse signals using a small number of linear measurements. In this work we investigate whether the measurement matrices used in compressed sensing can also be used to encrypt signals. Several papers allude to this possibility. Duarte et al. [6] refer to the compressed sensing measurements as “weakly encrypted” for an attacker without knowledge of the measurement matrix. Drori [7] states that “the encryption matrix can be viewed as a one-time pad that is completely secure.” In this paper we investigate the secrecy properties of compressed sensing measurements.

Joint compression and encryption could be useful in a number of applications where implementing an additional software layer for cryptography could be costly. For example, Akyildiz et al. [8] state that power consumption in sensor network nodes is a critical performance issue in many sensor network applications. Elimination of an additional protocol for encryption could be useful in this power-constrained scenario. To prevent privacy loss when databases are compromised, Draper et al. [9] propose lossy compression of biometric signals. Compressed sensing-based encryption could provide both signal compression and encryption guarantees, without the additional computational cost of a separate encryption protocol. Such computational savings could be significant when looking for biometric matches in a large database. Finally, high bandwidth sensors, such as video cameras, could jointly encrypt

and compress their measurements to reduce computational overhead in demanding applications such as video surveillance.

We now describe the organization of the paper. Section II provides an overview of compressed sensing, and introduces the secrecy definitions and model used in this paper. Section III discusses compressed sensing in the context of information-theoretic and computational definitions of secrecy. We will argue that compressed sensing measurements do not achieve information-theoretic secrecy. In Section IV, we present a theoretical result that demonstrates that compressed sensing measurements can achieve a computational notion of secrecy. This result provides a setting in which compressed sensing can be used to encrypt signals. Section V contains simulations that demonstrate the empirical performance of compressed sensing-based encryption. We conclude and discuss future work in Section VI.

II. BACKGROUND

A. Compressed Sensing

The work of Candes, Romberg, and Tao [1] and Donoho [3] showed that if a signal has a sparse representation in one basis then it can be recovered from a small number of projections onto a second basis that is incoherent with the first. We define sparsity and discuss incoherence below.

To define sparsity precisely, we introduce the following notation. Let Ψ be a matrix whose columns form an orthonormal basis. Define a K -sparse vector $\mathbf{x} \in \mathbb{R}^n$ as $\mathbf{x} = \Psi\theta$, where $\theta \in \mathbb{R}^N$ has K non-zero entries (i.e., it is K -sparse). We define Ω_K as the set of K indices over which the vector θ is non-zero.

A vector of measurements $\mathbf{y} = \Phi\mathbf{x}$, where Φ is an $M \times N$ matrix, is obtained by projecting the vector $\mathbf{x}$ onto a basis that is incoherent with Ψ . Roughly speaking, incoherence means that no basis vector in Ψ has a sparse representation in the basis specified by Φ . This notion is formalized in the compressed sensing literature [1–3]. We assume throughout that Φ is obtained by sampling independent and identically distributed (i.i.d.) Gaussian random variables with zero mean and variance $\frac{1}{M}$; Candes and Tao [2] showed that such a matrix is incoherent with high probability relative to any fixed basis Ψ . We

Y. Rachlin is with Accenture Technology Labs, 161 N Clark St., Chicago, IL 60601 yaron.rachlin@alumni.cmu.edu

D. Baron is with the Department of Electrical Engineering, Technion, Haifa 32000, Israel barondror@gmail.com

define Φ_{Ω_K} as the $M \times K$ measurement matrix obtained by selecting the K columns of Φ corresponding to the indices Ω_K .

For a K -sparse signal $\mathbf{x}$, only $K + 1$ projections of the signal onto the incoherent basis are required to reconstruct the signal with high probability [10, Theorem 2]. Unfortunately, this requires a combinatorial search, which is prohibitively complex. Candes et al. [1] and Donoho [3] proposed tractable recovery procedures based on linear programming, demonstrating the remarkable property that such procedures provide the same result as the combinatorial search as long as cK projections are used to reconstruct the signal (typically $c \approx 3$ or 4) [11, 12]. Conditions on the Φ matrix that enable tractable recovery methods can be stated in terms of the following property, introduced by Candes and Tao [11]. A Φ matrix is said to satisfy a restricted isometry property (RIP) of order K if there exists a $\delta_K \in (0, 1)$ such that,

$$(1 - \delta_K)\|\mathbf{x}\|_2^2 \leq \|\Phi\mathbf{x}\|_2^2 \leq (1 + \delta_K)\|\mathbf{x}\|_2^2 \quad (1)$$

holds for all $\mathbf{x}$ with sparsity K .

B. Secrecy Definitions and Model

To discuss the secrecy properties of compressed sensing measurements, we introduce the following conventions. A K -sparse message $\mathbf{x} \in \mathbb{R}^N$ is chosen by nature, possibly by sampling from some probability distribution. A key $i \in \{1, \dots, S\}$ corresponds to an $M \times N$ matrix Φ_i . In the model we consider in this paper, Alice wants to send a secret message to Bob. Alice chooses a key i (with uniform probability among the keys) and encrypts the message $\mathbf{x}$ using the Φ_i matrix via matrix multiplication $\mathbf{y} = \Phi_i\mathbf{x}$. Only the cryptogram $\mathbf{y}$ is transmitted to Bob, who knows what key that was used to encrypt the message. Given Φ_i and $\mathbf{y}$, the compressed sensing literature (discussed in Section II-A) provides conditions on $\mathbf{x}$ and Φ_i to enable the recovery of the message $\mathbf{x}$. We assume that given knowledge of the sparsity of $\mathbf{x}$, all Φ matrices obey these conditions, and therefore knowledge of the key enables Bob to recover the message.

Is the message that Alice sent to Bob secure? We assume that an eavesdropper, Eve, intercepts Alice's encrypted message $\mathbf{y}$, but does not know which key was used to encrypt the message. In this paper we study how difficult it is for Eve to recover $\mathbf{x}$ using only knowledge of $\mathbf{y}$, the sparsity of the message, and the set of keys and their corresponding Φ matrices.

An encryption system where both sender and receiver use the same key to respectively encrypt and decrypt the message is known as a secret key system [13]. In

compressed sensing, measurement matrices Φ can be generated randomly [2]. For such matrices, the secret key could therefore correspond to the seed of a pseudo-random number generator. Among other work, Diffie and Hellman [14] and Maurer and Wolf [15] address the secure exchange of secret keys. The attack we consider in this paper, where the eavesdropper observes only $\mathbf{y}$, is referred to as a ciphertext-only attack [13].

C. Related Work

McEliece [16] introduced a cryptographic system based on the difficulty of decoding linear codes with an unbounded number of errors. While both this work and our work consider linear encoders, in the McEliece system the error vector is the secret, while in our work the measurement matrix is the shared secret. In addition, our work considers the sparsity of the underlying signal being encrypted. Dwork et al. [17] apply ideas from compressed sensing to the problem of privacy in databases. Their paper demonstrates bounds on the fraction of errors a database query mechanism must introduce in order to prevent an adversary from reconstructing the database. The adversary chooses to design the measurement matrix by constructing queries. In our work, the measurement matrix is hidden from the adversary and no errors are artificially introduced.

III. NOTIONS OF SECRECY

How difficult is it for an eavesdropper to recover $\mathbf{x}$ without knowledge of the key used to encrypt the message? Discussion of the difficulty of breaking an encryption method can be broadly divided into computational and information theoretic approaches [18].

A. Perfect Secrecy

Information theoretic secrecy relies on the statistical properties of a system, and provides protection even in the face of a computationally unbounded adversary [18]. Shannon's work [19] pioneered this approach by introducing the idea of perfect secrecy. An encryption scheme achieves perfect secrecy if the probability of a message conditioned on the cryptogram is equal to the a priori probability of the message, $P(\mathbf{X} = \mathbf{x} | \mathbf{Y} = \mathbf{y}) = P(\mathbf{X} = \mathbf{x})$. Alternatively, this condition can be stated as $I(\mathbf{X}; \mathbf{Y}) = 0$. The following lemma demonstrates by contradiction that compressed sensing-based encryption does not achieve perfect secrecy.

Lemma 1: Let $\mathbf{X}$ be a message, $P_{\mathbf{X}}(\mathbf{x}) > 0 \quad \forall \mathbf{x} \in \mathbb{R}^n$, and Φ be an $M \times N$ measurement matrix. For $\mathbf{Y} = \Phi\mathbf{X}$,

$$I(\mathbf{Y}; \mathbf{X}) > 0,$$

and therefore perfect secrecy is not achieved.

Proof: $I(\mathbf{X}; \mathbf{Y}) = 0$ if and only if $\mathbf{X}$ and $\mathbf{Y}$ are independent [20, Theorem 8.6.1]. Since Φ is linear, $\mathbf{x} = 0$ implies that $\mathbf{y} = 0$. Therefore, $P_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y} = 0|\mathbf{X} = 0) = 1$. However, only $\mathbf{x}$ in the nullspace of Φ are mapped to $\mathbf{y} = 0$; by assumption, $P_{\mathbf{X}}(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathbb{R}^n$, and we conclude that $P_{\mathbf{Y}}(\mathbf{Y} = 0) < 1$. Therefore $P_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y} = 0|\mathbf{X} = 0) \neq P_{\mathbf{Y}}(\mathbf{Y} = 0)$, and $\mathbf{X}$ and $\mathbf{Y}$ are statistically dependent. ■

While the proof above depends on the fact that Φ is linear, other properties of the measurement matrix that are frequently used in compressed sensing can be used to prove that perfect secrecy is not achieved even when $\mathbf{X} \neq 0$. For example, assume Φ satisfies the restricted isometry property (1), and let $\mathbf{y} = \Phi\mathbf{x}$. Then,

$$\frac{\|\mathbf{y}\|_2^2}{(1 + \delta_K)} \leq \|\mathbf{x}\|_2^2. \quad (2)$$

For a fixed $\mathbf{y}$, choose $\mathbf{x}'$ such that $\|\mathbf{x}'\|_2^2 < \frac{\|\mathbf{y}\|_2^2}{(1 + \delta_K)}$. Using (2), we can state $P_{\mathbf{X}|\mathbf{Y}}(\mathbf{x}'|\mathbf{y}) = 0$. However, if we assume $P_{\mathbf{X}}(\mathbf{x}) > 0$ (e.g., a Gaussian distribution), then $P_{\mathbf{X}|\mathbf{Y}}(\mathbf{x}'|\mathbf{y}) \neq P_{\mathbf{X}}(\mathbf{x}')$, and $\mathbf{X}$ and $\mathbf{Y}$ are statistically dependent. The argument above relies on the fact that Φ that satisfy RIP roughly preserve norms, and therefore $\mathbf{y}$ provides information about the norm of the message $\mathbf{x}$. The arguments above do not preclude the possibility that $I(\mathbf{X}; \mathbf{Y})$ could be small; we leave the precise derivation of $I(\mathbf{X}; \mathbf{Y})$ for future research.

B. Computational Secrecy

Computation-based approaches such as public key cryptography are practical and widely used, but rely on open questions in complexity theory such as the difficulty of factoring [13]. In contrast to an information-theoretic notion of secrecy, the ciphertext in computational secrecy contains complete information about the message. However, extracting this information for an adversary without the appropriate key is equivalent to solving a computational problem that is assumed to be difficult (e.g., NP-hard). Thus, the statement of secrecy relies on assumptions about the difficulty of the computational problem and the computational resources available to the adversary. We will discuss the computational secrecy of compressed sensing in Section IV.

IV. SECRECY RESULT FOR COMPRESSED SENSING

How difficult is it for Eve to recover the message $\mathbf{x}$ using only $\mathbf{y}$, knowledge that $\mathbf{x}$ is K -sparse, and the set of keys and their corresponding Φ matrices? One possible approach is for Eve to methodically try all keys, and attempt to recover the signal $\mathbf{x}$, stopping when she thinks she has succeeded. Assume that Eve guesses

the measurement matrix Φ' . Then, she could attempt to recover using either ℓ_0 optimization,

$$\min_{\theta'} \|\theta'\|_0 \text{ subject to } \mathbf{y} = \Phi'\Psi\theta', \quad (3)$$

or using ℓ_1 optimization,

$$\min_{\theta'} \|\theta'\|_1 \text{ subject to } \mathbf{y} = \Phi'\Psi\theta'. \quad (4)$$

These two approaches and approximations to these optimization problems correspond to the majority of signal recovery methods for the setting of compressed sensing (i.e., strictly sparse signals and noiseless measurements) we consider in this paper. We emphasize that Eve may use other signal recovery methods than the ones described above. We will present results about Eve's ability to recover the secret message using (3) and (4) in Section IV-A. The implications of this analysis for the notions of secrecy introduced in Section III are discussed in Section IV-B.

A. Signal Recovery Using the Wrong Key

What can we state about the solution of either (3) or (4) when Eve guesses the wrong key k' ? We will show in Corollary 1 that when k' is different from the true key k (i.e., matrix Φ' is different than Φ), with probability one the attacker will recover an M -sparse solution instead of the original K -sparse signal. To establish the corollary, we will use Theorem 1, which proves that for a randomly generated Φ' matrix, all explanations of the measurements $\mathbf{y}$ are M -sparse.

Theorem 1: Let Φ and Φ' be two $M \times N$ matrices, randomly generated by sampling i.i.d. Gaussian random variables. For a K -sparse vector $\mathbf{x} = \Psi\theta$, let $\mathbf{y} = \Phi\mathbf{x}$. If $M \geq K + 1$, then all $\mathbf{x}' = \Psi\theta'$ such that $\mathbf{y} = \Phi'\mathbf{x}'$ satisfy $\|\theta'\|_0 = M$ with probability one over the set of Φ and Φ' matrices.

Proof: We begin by proving that there exists a unique solution for each M -set. Since Ψ is orthonormal, and the entries of Φ and Φ' are generated by sampling i.i.d. Gaussian random variables, the entries of the $M \times N$ matrices $\Phi\Psi$ and $\Phi'\Psi$ will also be Gaussian. Without loss of generality, we assume Ψ is the identity matrix, so $\mathbf{y} = \Phi\theta$.

A set of M columns of Φ' , indexed by Ω_M , are linearly independent with probability one over Φ' . Thus Φ'_{Ω_M} will have rank M , and matrix inversion can be used to uniquely determine M entries of θ that satisfy $\mathbf{y} = \Phi'\theta$. Therefore, with probability one, each set of M columns of Φ can be used to determine an $\mathbf{x}'$ that satisfies $\mathbf{y} = \Phi'\mathbf{x}'$.

Next, we prove lack of a T -sparse solution, where $T < M$. Let Ω_T denote the indices of the non-zero entries of a T -sparse vector θ' . The matrix Φ'_{Ω_T} has rank T with probability one over Φ' . The indices of the K non-zero entries of the vector θ are Ω_K , and we note that Φ_{Ω_K} has rank K with probability one. Denote by $\text{colspan}(A)$ the vector space spanned by columns of A .

We will show that if $M \geq K + 1$ then $y \notin \text{colspan}(\Phi'_{\Omega_T})$ with probability one over Φ and Φ' . To show this, we analyze the concatenated matrix $[\Phi_{\Omega_K} \Phi'_{\Omega_T}]$ in two different cases. First, consider $K + T > M$. In this case the concatenated matrix has rank M with probability one over Φ and Φ' , because the rows consist of i.i.d. Gaussian random variables. Therefore, $\text{colspan}(\Phi_{\Omega_K}) \cap \text{colspan}(\Phi'_{\Omega_T})$ has dimension $K + T - M$. Since by definition $T < M$, the intersection has dimension less than K . The measurements can only be expressed as a sum of the columns of Φ'_{Ω_T} if they lie in this intersection. However, with probability one over the Φ matrices, these measurements lie in a K -dimensional space. Since there are a finite number of subsets of T indices, the probability that a T -sparse vector x' will satisfy $y = \Phi'x'$ is zero over Φ .¹

We now consider $K + T \leq M$. In this case, the concatenated matrix $[\Phi_{\Omega_K} \Phi'_{\Omega_T}]$ has rank $K + T$ with probability one since the columns are generated as i.i.d. Gaussian random variables. Therefore $\text{colspan}(\Phi_{\Omega_K}) \cap \text{colspan}(\Phi'_{\Omega_T})$ has dimension 0. In this case y will not be in the column span of Φ'_{Ω_T} . ■

Corollary 1 below is a simple consequence of Theorem 1. This corollary demonstrates that with probability one, when Eve uses the wrong key, then the recovered solution will be an M -sparse signal, instead of the true K -sparse signal. While our corollary only holds for the recovery methods (3) and (4), the idea that recovery using the wrong key is difficult should be intuitively clear, and will be strengthened in future work.

Corollary 1: Let Φ and Φ' be $M \times N$ matrices with entries generated by sampling i.i.d. Gaussian random variables. Let x be K -sparse and $y = \Phi x$. When $M \geq K + 1$, the ℓ_0 optimization (3) and the ℓ_1 optimization (4) will yield an M -sparse solution with probability one.

Proof: Both (3) and (4) describe optimization problems over the set of θ' that satisfy $y = \Phi'\Psi\theta'$. According to Theorem 1, with probability one all vectors θ' that satisfy this constraint are M -sparse. Consequently, the ℓ_0 optimization problem (3) and the ℓ_1 optimization problem (4) can only yield M -sparse solutions. ■

¹Our statements that a set occurs with probability zero mean that a set has Lebesgue measure zero. Since the distributions we consider are continuous, the two statements are equivalent.

B. Secrecy Implications

How do our results address the different notions of secrecy introduced in Section III? Assume that Eve attacks using either (3) or (4). Then with regard to the computational notion of secrecy, she needs to methodically evaluate keys to find the Φ that recovers a K -sparse solution instead of an M -sparse solution. The amount of computation required to find the correct key is proportional to the number of keys S . In practice, $S = 2^{64}$ could be accomplished by sharing a sufficiently large random seed, making a methodical evaluation of all keys difficult. Once Eve recovers a K -sparse vector using the correct key, she knows with probability one that it is the correct key. The encryption is computationally hard to crack for a large number of keys, but a computationally unbounded adversary can easily infer that the correct K -sparse signal has been recovered. The ability to identify the correct solution is a result of the dependence between x and y , and demonstrates the lack of perfect secrecy in compressed sensing-based encryption. Thus in the case of noiseless measurements and strictly sparse signals, and an attacker attempting to recover the signal using either (3) or (4), compressed sensing measurements provide computational secrecy.

To demonstrate computational secrecy in the face of all feasible attacks, it is insufficient to show that a particular set of attacks will be computationally difficult. Would Eve be more successful if she employed different attacks than the ones analyzed in this paper? While this is a topic for future work, we present an argument which shows that any attack that *precisely* recovers x reduces to solving an NP-hard problem. Let $M > K$, and consider an algorithm that implements the following function,

$$x = A(y, \Phi_1, \dots, \Phi_S, K). \quad (5)$$

For example, such an algorithm could methodically try to reconstruct using all Φ matrices until a K -sparse reconstruction is found, and then outputs the corresponding x . Could such an algorithm run in polynomial time? If it would, then its output would be a solution to ℓ_0 minimization (3). However, as Candes et al. [11] point out, ℓ_0 reconstruction is NP-hard. The algorithm (5) would therefore demonstrate that $P = NP$, and would require significant new insight. This argument demonstrates that while other attacks are feasible, a successful attack that recovers x would, barring significant advances, not run in polynomial time.

V. SIMULATIONS

To provide a numerical demonstration of our work, we generated an ensemble of 100 matrices and 100 messages for each signal length N . Each message was

Fig. 1. Statistics of attacker reconstruction error in a simulated attack, as a function of signal length.

a spike signal, consisting of zeros except for spikes of magnitude one. The number of spikes was $K = 0.1N$ and the number of measurements was $M = 4K$. We encrypted each message using its corresponding key, and attempted to recover the encrypted message using the other 99 keys. For each message length N , Figure 1 shows the result of 9900 attacks. The lines in the graph correspond to the 90th percentile, median, 10th percentile, and lowest attacker reconstruction error, where error is measured in terms of the squared ℓ_2 norm of the reconstruction. The reconstruction was performed using LIMagic [21]. From the graph, it is apparent that Eve experiences significant reconstruction error that increases linearly with N . At the same time, Bob reconstructs with low error (with appropriate settings, up to floating point precision). For example, for $N = 269$ the average squared ℓ_2 reconstruction error (over 100 messages) when using the correct key was 9.6×10^{-5} .

VI. CONCLUSIONS AND FUTURE WORK

We presented an analysis of the secrecy properties of noiseless compressed sensing measurements of strictly sparse signals. We proved that such compressed sensing measurements do not achieve perfect secrecy. Instead, we demonstrated a computational notion of secrecy by showing that recovery using the wrong key will produce an incorrect signal with probability one for all keys but the one used to originally encrypt the signal. This analysis shows a setting in which compressed sensing measurements could be considered to be encrypted. Future work includes analysis of the secrecy of noisy measurements, and of compressible, as opposed to strictly sparse, signals.

Acknowledgements: We would like to thank Yoram Bresler, Rayid Ghani, Prakash Ishwar, Olgica Milenkovic, and Pierre Moulin for helpful comments.

REFERENCES

- [1] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52(2), pp. 489–509, 2006.
- [2] E. Candes and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52(12), pp. 5406–5425, 2006.
- [3] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52(4), pp. 1289–1306, 2006.
- [4] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, vol. 24(4), 2007.
- [5] E. Candes and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25(2), 2008.
- [6] M. F. Duarte, S. Sarvotham, M. B. Wakin, D. Baron, and R. G. Baraniuk, "Joint sparsity models for distributed compressed sensing," in *Online Proceedings of the Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, November 2005.
- [7] I. Drori, "Compressed video sensing," in *BMVA Symposium on 3D Video - Analysis, Display, and Applications*, 2008.
- [8] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, 2002.
- [9] S. C. Draper, A. Khisti, E. Martinian, A. Vetro, and J. S. Yedidia, "Using distributed source coding to secure fingerprint biometrics," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2007 (ICASSP)*, 2007.
- [10] D. Baron, M. B. Wakin, M. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed compressed sensing," submitted.
- [11] E. J. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51(12), pp. 4203–4215, December 2005.
- [12] D. Donoho and J. Tanner, "Neighborliness of randomly projected simplices in high dimensions," *Proc. National Academy of Sciences*, vol. 102, no. 27, pp. 9452–457, 2005.
- [13] R. Oppliger, *Contemporary Cryptography*. Artech House, 2005.
- [14] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. IT-22, no. 6, pp. 644–654, 1976.
- [15] U. Maurer and S. Wolf, "Information-theoretic key agreement: From weak to strong secrecy for free," in *Advances in Cryptology — EUROCRYPT 2000*, ser. Lecture Notes in Computer Science, B. Preneel, Ed., vol. 1807. Springer-Verlag, May 2000, pp. 351–368.
- [16] R. J. McEliece, "A public-key cryptosystem based on algebraic coding theory," Jet Propulsion Laboratory, Tech. Rep. DSN Progress Report 42-44, 1978.
- [17] C. Dwork, F. McSherry, and K. Talwar, "The price of privacy and the limits of LP decoding," in *Symp. on Theory of Computing (STOC)*, June 2007.
- [18] U. Maurer, "The role of information theory in cryptography," in *Proc. of 4th IMA Conference on Cryptography and Coding*, P. Farrell, Ed. The Institute of Mathematics and its Applications, Southend-on-Sea, England, Dec. 1993, pp. 49–71.
- [19] C. E. Shannon, "Communication theory of secrecy systems," *Bell System Technical Journal*, vol. 28(4), pp. 656–715, October 1949.
- [20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, 2006.
- [21] E. Candes and J. Romberg, "L1-magic," <http://www.l1-magic.org/>, 2007.